

09/14/00
jc9044U.S. PTO

09-18-00

EK287384749US
9/14/00

A

CERTIFICATE OF MAILING BY "EXPRESS MAIL" UNDER 37 CFR § 1.10

Express Mail® mailing label number _____

Date of Mailing: _____

I hereby certify that the documents indicated below are being deposited with the United States Postal Service under 37 CFR 1.10 on the date indicated above and are addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231, and mailed on the above Date of Mailing with the above "Express Mail" mailing label number _____

(Typed or printed name of person mailing paper or fee) _____ SIGNATURE of person mailing paper or fee _____

BOX PATENT APPLICATION
ASSISTANT COMMISSIONER FOR PATENTS
WASHINGTON, D. C. 20231

DOCKET NUMBER: AUS9-2000-0476-US1

Sir:

Transmitted herewith for filing is the Patent Application of:

Inventors: Deanna L. Quigg Brown et al.

For: DETECTION OF DEAD GATEWAYS

Enclosed are:

- ☒ Patent Specification and Declaration
- ☒ 4 sheets of drawing(s)
- ☒ An assignment of the invention to International Business Machines Corporation (includes Recordation Form Cover Sheet).
- ☐ A certified copy of a application.
- ☐ An associate power of attorney
- ☒ Information Disclosure Statement Certificate

The filing fee has been calculated as shown below:

For	Number Filed	Number Extra	Rate	Fee
Basic Fee				\$ 690.00
Total Claims	75 - 20	55	x 18 =	\$ 990.00
Indep. Claims	3 - 3	0	x 78 =	\$ - 0 -
<input type="checkbox"/> MULTIPLE DEPENDENT CLAIM(S) PRESENTED			+ 260 =	\$ - 0 -
TOTAL				\$ 1680.00

- ☒ Please charge my Deposit Account No. 09-0447 in the amount of \$1680.00. **A duplicate copy of this sheet is enclosed.**
- ☒ The Assistant Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to Deposit Account No. 09-0447. **A duplicate copy of this sheet is enclosed.**
 - ☒ Any additional filing fees required under 37 CFR §1.16
 - ☒ Any patent application processing fees under 37 CFR §1.17.

Respectfully submitted,

By: _____

Volol Emile
Registration No. 39,969
IBM Corporation
Intellectual Property Law Dept.
Internal Zip 9444
11400 Burnet Road
Austin, Texas 78758
Telephone: (512) 823-1005

::ODMA\PCDOCS\AUSTIN_1\147214\1
207:7047-P366US

jc825 U.S. PTO
09/661275

09/14/00

09-18-00

DETECTION OF DEAD GATEWAYS

TECHNICAL FIELD

5 The present invention relates to the field of digital packet telecommunications, and more particularly to a TCP/IP host configured to detect a dead gateway.

BACKGROUND INFORMATION

10 Digital packets of information may be communicated between an organization's internal network and the public Internet. The Internet is a collection of networks, including Arpanet, NSFnet, regional networks such as NY sernet, local networks at a number of university and research institutions, and a number of military networks. Many organizations use Internet gateways, implemented by computer workstations executing software, to provide a connection between the internal network and the Internet. These Internet gateways can include firewall protection for the internal network in addition to providing connectivity between the internal network and the Internet. Where an organization uses more than one Internet gateway, each user on the organization's internal network can, for example, be assigned to a designated Internet gateway. The users of the internal network are thereby provided access to the Internet through their designated Internet gateways.

20 Information is typically transmitted over the Internet using Transmission Control Protocol/ Internet Protocol (TCP/IP) protocol suite. Figure 1 illustrates the constituents of the TCP/IP protocol suite 100. The base layer of the TCP/IP protocol suite is the physical layer 110, which defines the mechanical, electrical, functional and procedural standards for the physical transmission of data over communications media. The
25 physical layer may comprise electrical, mechanical or functional standards such as whether a network is packet switching or frame-switching or whether a network is based on a Carrier Sense Multiple Access/Collision Detection (CSMA/CD).

Overlying the physical layer 110 is the data link layer 120. Data link layer 120 provides the function and protocols to transfer data, e.g., transfer data between clients of different networks, and to detect errors that may occur at the physical layer. Operating modes at the data link layer 120 comprise such standard network topologies as IEEE 802.3 Ethernet, IEEE 802.5 Token Ring, ITU X.25, or serial (SLIP) protocols.

Network layer protocols 130 overlay the data link layer 120 and provide the means for establishing connections between networks. The standards of network layer protocols 130 provide operational control procedures for internetworking communications and routing information through multiple heterogenous networks. Examples of network layer protocols are the Internet Protocol (IP) and the Internet Control Message Protocol (ICMP). The Address Resolution Protocol (ARP) is used to correlate an Internet address and a Media Access Address (MAC) for a particular host. The Routing Information Protocol (RIP) is a dynamic routing protocol for passing routing information between hosts on networks. The Internet Control Message Protocol ICMP is an internal protocol for passing control messages between hosts on various networks. ICMP messages provide feedback about events in the network environment or can help determine if a path exists to a particular host in the network environment. The latter is commonly referred to as a "ping". The Internet Protocol (IP) provides the basic mechanism for routing packets of information in the Internet. IP is a non-reliable communication protocol. It provides a "best efforts" delivery service and does not commit network resources to a particular transaction, nor does it perform retransmissions or give acknowledgments.

The transport layer 140 provides end-to-end transport services across multiple heterogenous networks. The User Datagram Protocol (UDP) provides a connectionless, datagram oriented service which provides a non-reliable delivery mechanism for streams of information. The Transmission Control Protocol (TCP) provides a reliable session-based service for delivery of sequenced packets of information across the Internet. TCP provides a connection oriented reliable mechanism for information delivery.

5 The application layer 150 provides a list of network applications and utilities, a few of which are illustrated in Figure 1. For example, File Transfer Protocol (FTP) is a standard TCP/IP protocol for transferring files from one computer to another. FTP clients establish sessions through TCP connections with FTP servers in order to obtain files. Telnet is a standard TCP/IP protocol for remote terminal connection. The Hypertext Transfer Protocol (HTTP) facilitates the transfer of data objects across networks via a system of uniform resource indicators (URL). The Simple Network Management Protocol (SNMP) is a standard for managing TCP/IP networks. A Remote Procedure Call (RPC) is a programming interface which enables programs to invoke remote functions on server machines.

10

Internet gateways generally use routing protocols, e.g., Open Shortest Path First (OSPF), to adjust to changes in the network. That is, routing of the digital packets of information may continue despite the fact that some routers are inoperative. However, Internet hosts generally do not participate in routing protocols. If a host's first-hop gateway, i.e., first gateway to the Internet host, becomes inoperative, the host will not adjust and will continue to transmit digital packets of information through a gateway that is inoperative. Hence, the host is effectively transmitting digital packets of information down a "black hole." For this reason, RFC1122, the document which governs the routing of digital packets of information across the Internet states:

20 "The IP layer must be able to detect the failure of a 'next-hop' gateway that is listed in its route cache and to choose an alternative gateway."

One prior art solution to the host continuing to transmit digital packets of information to the dead gateway is to simply try another route if the TCP connection is failing. Since TCP provides feedback to a host if the transmission is failing, the host may be able to select another route through a different router or gateway to send the TCP traffic. However, non-TCP traffic involves protocols that do not provide feedback to a host if the transmission is failing. Since address resolution protocol (ARP) entries which obtain a data link layer address from the corresponding network layer address (the

25

address that IP uses) in the host are cached for a certain duration, e.g., twenty minutes, non-TCP traffic will fail for as long as the ARP entry is cached despite the fact that there was an alternative gateway. That is, non-TCP traffic will continue to be transmitted to the first-hop gateway, i.e., gateway that is one hop from the host sending the data, that is inoperative.

It would therefore be desirable to allow a host to detect a first-hop dead gateway, i.e., gateway that is inoperative, and be able to send non-TCP data through an alternative route upon a subsequent non-TCP communication. It would further be desirable to allow the host to mark the first-hop dead gateway to a lower priority if it remains inoperative upon a subsequent new communication.

SUMMARY

5 The problems outlined above may at least in part be solved in some embodiments by an application of a sender host sending a TCP packet of data through a first gateway to a receiver host. The first gateway is a first-hop away from the sender host. If TCP fails to receive an acknowledgment of received data from the receiver host, the ARP entry associated with the first-hop gateway is deleted in the sender host. When the application or a new application of the sender host establishes a new communication with a route through the gateway whose ARP entry was deleted, an ARP request is sent to that gateway by the sender host. If the sender host receives a response from the gateway whose ARP entry was previously deleted, then that gateway is operative. A TCP or non-TCP packet of data may then be sent by the application or new application of the sender host to the receiver host through the first-hop gateway whose ARP entry was previously deleted if the new communication established was TCP or non-TCP, respectively. If the sender host does not receive a response from the gateway whose ARP entry was previously deleted, then that gateway is inoperative, i.e., dead. The sender host then selects an alternative path through an alternative first-hop gateway in a routing table in the sender host. A TCP or non-TCP packet of data may then be sent to the receiver host by the application or new application of the sender host through the alternative first-hop gateway if the new communication was TCP or non-TCP, respectively.

25 In one embodiment, a method for detecting a first-hop dead gateway comprises the step of sending a TCP packet of data from an application of a sender host to a receiver host through a first gateway, where the first gateway is a first-hop away from the sender host. The method further comprises the step of TCP failing to receive an acknowledgment of received data from the receiver host. The method further comprises the step of deleting an ARP entry associated with the first gateway in the sender host. The method further comprises the step of selecting an alternative path to send the TCP

packet of data from the sender host to the receiver host through a second gateway in a routing table in the sender host, where the second gateway is a first-hop away from the sender host. The method further comprises the step of establishing a new communication using the first gateway by the application or new application of the sender host. The method further comprises the step of sending an ARP request to the first gateway by the sender host. If the sender host receives a response from the first gateway, then the method further comprises the step of sending a TCP or non-TCP packet of data using the first gateway if the new communication was a TCP or non-TCP communication, respectively. If the sender host does not receive a response from the first gateway, then the method further comprises the step of selecting an alternative path through an alternative first-hop gateway in the routing table of the sender host. The application or new application of the sender host then sends the TCP or non-TCP packet of data using the alternative gateway if the new communication was a TCP or non-TCP communication, respectively.

In another embodiment of the present invention, if the sender host does not receive a response from the first gateway, i.e., gateway whose ARP entry was deleted, then all routes using that gateway are marked to a lower priority level in the routing table in the sender host. In another embodiment of the present invention, the routes through the gateway that were marked to a lower priority level may return to their original priority level upon a duration of time, e.g., five minutes.

The foregoing has outlined rather broadly the features and technical advantages of the present invention in order that the detailed description of the invention that follows may be better understood. Additional features and advantages of the invention will be described hereinafter which form the subject of the claims of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

A better understanding of the present invention can be obtained when the following detailed description is considered in conjunction with the following drawings, in which:

Figure 1 depicts a relationship diagram of the layers of the TCP/IP protocol suite;

Figure 2 illustrates transferring digital packets of information between a host on one network to a host on another network/Internet;

Figure 3 illustrates a host configured in accordance with the present invention; and

Figure 4 is a flowchart depicting a method of a host detecting a first-hop dead gateway.

DETAILED DESCRIPTION

5 The present invention comprises a method, computer program product and system for detecting a first-hop dead gateway. In one embodiment of the present invention, an application of a sender host sends a TCP packet of data through a first gateway to a receiver host. The first gateway is a first-hop away from the sender host. If TCP fails to receive an acknowledgment of received data from the receiver host, the ARP entry associated with the first-hop gateway is deleted in the sender host. When the application or new application of the sender host establishes a new communication with a route through the gateway whose ARP entry was deleted, an ARP request is sent to that gateway by the sender host. If the sender host receives a response from the gateway whose ARP entry was previously deleted, then that gateway is operative. A TCP or non-TCP packet of data may then be sent by the application or new application of the sender host to the receiver host through the first-hop gateway whose ARP entry was previously deleted if the new communication established was TCP or non-TCP, respectively. If the sender host does not receive a response from the gateway whose ARP entry was previously deleted, then that gateway is inoperative, i.e., dead. The sender host then selects an alternative path through an alternative first-hop gateway in a routing table in the sender host. A TCP or non-TCP packet of data may then be sent to the receiver host by the application or new application of the sender host through the alternative first-hop gateway if the new communication was TCP or non-TCP, respectively. In another embodiment of the present invention, all routes using the dead gateway are marked to a lower priority level in the routing table in the sender host.

25 Figure 2 - Embodiment of a Network connected to another Network through a plurality of Gateways

Figure 2 illustrates an embodiment of the present invention where network 230A is connected to another network 230B, e.g., Internet, through gateways 220A-B. A

plurality of hosts 210A-C is coupled to network 230A. A plurality of hosts 210D-F is coupled to network 230B. Hosts 210A-F may collectively or individually be referred to as hosts 210 or host 210, respectfully. Gateways 220A-B may collectively or individually be referred to as gateways 220 or gateway 220, respectfully. Networks 230A-B may collectively or individually be referred to as networks 230 or network 230, respectfully. It is noted that any number of hosts 210 may be coupled to networks 230. It is further noted that one of the networks 230 may represent the Internet which is a collection of networks.

Referring to Figure 1, each of the plurality of hosts 210A-C of network 230A is configured to send a packet a data, e.g., TCP and non-TCP data, to any of the plurality of hosts of network 230B through one of the plurality of gateways 220A-B. Each of the plurality of hosts 210D-F of network 230B is configured to send a packet a data, e.g., TCP and non-TCP data, to any of the plurality of hosts of network 230A through one of the plurality of gateways 220A-B. For example, an application running on host 210A may send a TCP or non-TCP packet of data to host 210D though gateway 220A. As stated in the Background Information section, when a first-hop gateway, i.e., gateway that is one hop from the host sending the data, becomes inoperative, the host will not adjust for non-TCP data and will continue to transmit non-TCP digital packets of information to the gateway that is inoperative. It would therefore be desirable to allow a host to detect a first-hop dead gateway, i.e., gateway that is inoperative, and be able to send non-TCP data through an alternative route upon a subsequent non-TCP communication. It would further be desirable to allow the host to mark the first-hop dead gateway to a lower priority if it remains inoperative upon a subsequent new communication. A detailed description of host 210 that is configured to detect a first-hop dead gateway and be able to send non-TCP data through an alternative route upon a subsequent non-TCP communication is described below in Figure 3. A method of host 210 detecting a first-hop dead gateway is provided in Figure 4.

Figure 3- Embodiment of a Host

Figure 3 illustrates a typical hardware configuration of host 210 which is representative of a hardware environment for practicing the present invention. Host 210 has a central processing unit (CPU) 310, such as a conventional microprocessor, coupled to various other components by system bus 312. Read only memory (ROM) 316 is coupled to system bus 312 and includes a basic input/output system ("BIOS") that controls certain basic functions of host 210. Random access memory (RAM) 314, I/O adapter 318, and communications adapter 334 are also coupled to system bus 312. I/O adapter 318 may be a small computer system interface ("SCSI") adapter that communicates with disk units 320 and tape drives 340. Communications adapter 334 interconnects bus 312 with an outside network enabling host 210 to communicate with other such systems. Input/Output devices are also connected to system bus 312 via a user interface adapter 322 and a display adapter 336. A display monitor 338 is connected to system bus 312 by display adapter 336. In this manner, a user is capable of inputting to host 210 through a keyboard 324 or a mouse 326 and receiving output from host 210 via display 238.

Preferred implementations of the invention include implementations as a host programmed to execute the method or methods described herein, and as a computer program product. According to the host implementations, sets of instructions for executing the method or methods are resident in the random access memory 314 of one or more hosts configured generally as described above. Until required by the host, the set of instructions may be stored as a computer program product in another computer memory, for example, in disk drive 320 (which may include a removable memory such as an optical disk or floppy disk for eventual use in disk drive 320). Furthermore, the computer program product can also be stored at another host and transmitted when desired to the user's work station by a network or by an external network such as the Internet. One skilled in the art would appreciate that the physical storage of the sets of instructions physically changes the medium upon which it is stored so that the medium

carries computer readable information. The change may be electrical, magnetic, chemical or some other physical change.

Figure 4 - Flowchart of a Method for Detecting a First-Hop Dead Gateway

Figure 4 illustrates a flowchart of one embodiment of the present invention of a method 400 for detecting a first-hop dead gateway and sending non-TCP data through an alternative route upon a subsequent non-TCP communication. As stated above, a host of one network, e.g., 210A, is configured to send a packet of data, e.g., TCP or non-TCP data, to a host of another network, e.g., 210D, through a gateway 220 as illustrated in Figure 2. The host that sends the packet of data is commonly referred to as a sender host. The host that receives the packet of data is commonly referred to as a receiver host.

In step 405, an application operating on the sender host, e.g., 210A, sends a TCP packet of data to the receiver host, e.g., 210D, through a gateway, e.g., 220A, in a TCP connection. A determination is then made in step 410 as to whether TCP received an acknowledgment of received data from the receiver host as described below. As stated in the Background Information section, TCP provides a connection oriented reliable mechanism for information delivery. One method employed by TCP to provide reliability is to expect each end, i.e., sender host and receiver host, to acknowledge the data received from the other end. At times though, data segments and acknowledgments may be lost. To solve this problem, TCP sets a time-out when it sends data so that the data unacknowledged is retransmitted. Hence, if TCP does not receive an acknowledgment of received data from the receiver host, then TCP resends the TCP data packet to the receiver host, e.g., 210D, through the same gateway, e.g., 220A. In one embodiment, the process may be repeated for a certain consecutive number of times or until TCP receives an acknowledgment of received data from the receiver host. If TCP receives an acknowledgment of received data from the receiver host, then method 400 ends at step 415.

If TCP did not receive an acknowledgment of received data from the receiver host, then an Address Resolution Protocol (ARP) entry associated with the first-hop

gateway is deleted in the sender host in step 420. TCP may not receive an acknowledgment of received data for a number of reasons including: lost data segments or lost acknowledgments, the TCP data packet may be sent to a non-existent address, the TCP data packet may be dropped by an intermediary router or gateway, the route may be administratively prohibited, the first-hop gateway is dead, etc. Since there are other reasons for the TCP to not receive an acknowledgment of received data from the receiver other than an inoperative first-hop gateway, an ARP entry associated with the first-hop gateway is deleted in the sender host after the TCP has not received an acknowledgment of received data from the receiver host in step 420. In one embodiment, the ARP entry associated with the first-hop gateway is deleted in the sender host after the TCP has not received an acknowledgment of received data from the receiver host for a consecutive number of times. By deleting the ARP entry, the sender host may be able to determine if the first-hop gateway was indeed inoperative upon a new communication as described in steps 430 and 435.

Upon deletion of the ARP entry, an alternative route, i.e., path, is selected by the sender host to send the TCP packet of data to the receiver host through an alternative first-hop gateway in a routing table in the sender host in step 425. The alternative path is selected prior to the TCP connection timing out. For example, sender host, e.g., 210A, may now attempt to send the TCP packet of data through a different first-hop gateway, e.g., 220B, by selecting the route which includes the new gateway, e.g., 220B, in the routing table. The routing table is typically located in the network layer, specifically the IP layer of the TCP/IP protocol suite. In one embodiment, the TCP/IP protocol suite may be located in the kernel memory on the sender host.

In step 430, the application or a new application of the sender host attempts to establish a new communication, TCP or non-TCP, using the first-hop gateway whose ARP entry was deleted in step 420. When the application or new application of the sender host attempts to establish a new communication, TCP or non-TCP, using the gateway whose ARP entry was deleted in step 420, the sender host sends an ARP request to that gateway in step 435. The ARP request is used to correlate an Internet Protocol

address and a Media Access Address (MAC) for that gateway. A determination is then made in step 440 as to whether the sender host receives a response from that first-hop gateway. If that first-hop gateway is operative, then the sender host receives a response from that gateway. A TCP or non-TCP packet of data may then be sent by the application or new application of the sender host to the receiver host through the first-hop gateway whose ARP entry was deleted in step 420 if the new communication was TCP or non-TCP, respectively in step 445.

If the sender host fails to receive a response from the first-hop gateway whose ARP entry was deleted, then a new route, i.e., path, is selected by the sender host in a routing table in the sender host in step 450. If the gateway whose ARP entry was deleted is inoperative, i.e., dead, then the sender host fails to receive a response from that gateway. Subsequently, an alternative route, i.e., path, is selected in the routing table in the sender host where the route is through an alternative first-hop gateway than the dead gateway. In one embodiment, an ARP request may be sent a consecutive number of times by the sender host to the gateway whose ARP entry was deleted if the sender host fails to receive a response from the gateway whose ARP entry was deleted. A TCP or non-TCP packet of data may then be sent to the receiver host by the application or new application of the sender host through the alternative first-hop gateway if the new communication was TCP or non-TCP, respectively.

If the sender host does not receive a response from the gateway whose ARP entry was deleted in step 420, then all routes using that gateway are marked to a lower priority level in the routing table in the sender host in step 455. In one embodiment, the routes through the gateway that were marked to a lower priority level may return to their original priority level upon a duration of time, e.g., five minutes.

Although the method, computer program product and system of the present invention are described in connection with several embodiments, it is not intended to be limited to the specific forms set forth herein, but on the contrary, it is intended to cover such alternatives, modifications, and equivalents, as can be reasonably included within

the spirit and scope of the invention as defined by the appended claims. It is noted that the headings are used only for organizational purposes and not meant to limit the scope of the description or claims.

CLAIMS:

1. A method for selecting an alternative path upon detection of a dead gateway comprising the steps of:

 sending a first TCP packet of data from an application of a sender host to a receiver host through a first gateway;

 failing to receive an acknowledgment of received data from said receiver host;

 deleting an ARP entry associated with said first gateway in said sender host; and

 selecting an alternative path to send said first TCP packet of data from said sender host to said receiver host through a second gateway in a routing table in said sender host.

2. The method as recited in claim 1, wherein said first gateway is a first-hop away from said sender host, wherein said second gateway is a first-hop away from said sender host.

3. The method as recited in claim 1, wherein said first TCP packet of data is sent from said sender host to said receiver host through said first gateway a consecutive number of times, wherein said sender host fails to receive an acknowledgment of received data from said receiver host each time said first TCP packet of data is sent to said receiver host.

4. The method as recited in claim 1, wherein said routing table in said sender host is located in a network layer of a TCP/IP protocol suite.

5. The method as recited in claim 4, wherein said network layer is an IP layer.

6. The method as recited in claim 1 further comprising the step of:
 establishing a new communication using said first gateway by said application of said sender host.

1 7. The method as recited in claim 1 further comprising the step of:
2 establishing a new communication using said first gateway by a new application
3 of said sender host

1 8. The method as recited in claim 6 further comprising the step of:
2 sending an ARP request to said first gateway by said sender host.

1 9. The method as recited in claim 8, wherein said ARP request is sent from said
2 sender host to said first gateway a consecutive number of times.

1 10. The method as recited in claim 8, wherein if said sender host receives a response
2 from said first gateway then said first gateway is operative, wherein if said sender host
3 receives a response from said first gateway then said application of said sender host
4 sends a second TCP packet of data using said first gateway if said new communication
5 is a TCP communication.

1 11. The method as recited in claim 8, wherein if said sender host receives a response
2 from said first gateway then said first gateway is operative, wherein if said sender host
3 receives a response from said first gateway then said application of said sender host
4 sends a non-TCP packet of data using said first gateway if said new communication is
5 a non-TCP communication.

1 12. The method as recited in claim 8, wherein if said sender host does not receive a
2 response from said first gateway then said first gateway is inoperative, wherein if said
3 sender host does not receives a response from said first gateway then the method further
4 comprises the step of:
5 selecting an alternative path through an alterative gateway than said first gateway
6 in said routing table in said sender host.

1 13. The method as recited in claim 12, where said application of said sender host
2 sends a non-TCP packet of data using said alternative gateway if said new
3 communication is a non-TCP communication.

1 14. The method as recited in claim 12, where said application of said sender host
2 sends a second TCP packet of data using said alternative gateway if said new
3 communication is a TCP communication.

1 15. The method as recited in claim 12 further comprising the step of:
2 marking all routes that use said first gateway to a lower priority level from an
3 original priority level in said routing table in said sender host.

1 16. The method as recited in claim 15, wherein said all routes that use said first
2 gateway return to their original priority level after a duration of time.

1 17. The method as recited in claim 7 further comprising the step of:
2 sending an ARP request to said first gateway by said sender host.

1 18. The method as recited in claim 17, wherein said ARP request is sent from said
2 sender host to said first gateway a consecutive number of times.

1 19. The method as recited in claim 17, wherein if said sender host receives a response
2 from said first gateway then said first gateway is operative, wherein if said sender host
3 receives a response from said first gateway then said new application of said sender host
4 sends a second TCP packet of data using said first gateway if said new communication
5 is a TCP communication.

1 20. The method as recited in claim 17, wherein if said sender host receives a response
2 from said first gateway then said first gateway is operative, wherein if said sender host
3 receives a response from said first gateway then said new application of said sender host

4 sends a non-TCP packet of data using said first gateway if said new communication is
5 a non-TCP communication.

1 21. The method as recited in claim 17, wherein if said sender host does not receive
2 a response from said first gateway then said first gateway is inoperative, wherein if said
3 sender host does not receives a response from said first gateway then the method further
4 comprises the step of:

5 selecting an alternative path through an alterative gateway than said first gateway
6 in said routing table in said sender host.

1 22. The method as recited in claim 21, where said new application of said sender host
2 sends a non-TCP packet of data using said alternative gateway if said new
3 communication is a non-TCP communication.

23. The method as recited in claim 21, where said new application of said sender host
sends a second TCP packet of data using said alternative gateway if said new
communication is a TCP communication.

24. The method as recited in claim 21 further comprising the step of:
marking all routes that use said first gateway to a lower priority level from an
original priority level in said routing table in said sender host.

1 25. The method as recited in claim 24, wherein said all routes that use said first
2 gateway return to their original priority level after a duration of time.

1 26. A system for selecting an alternative path upon detection of a dead gateway
2 comprising:

3 a first and second network connected together via a first gateway;
4 at least one host electrically coupled to said first network, wherein one of said at
5 least one host coupled to said first network is a sender host;

6 at least one host electrically coupled to said second network, wherein one of said
7 at least one host coupled to said second network is a receiver host;

8 wherein said sender host further comprises:

9 circuitry for sending a first TCP packet of data from an application of said
10 sender host to said receiver host through said first gateway, wherein said sender host fails
11 to receive an acknowledgment of received data from said receiver host;

12 circuitry for deleting an ARP entry associated with said first gateway in
13 said sender host; and

14 circuitry for selecting an alternative path to send said first TCP packet of
15 data to said receiver host through a second gateway in a routing table in said sender host.

16 27. The system as recited in claim 26, wherein said first gateway is a first-hop away
17 from said sender host, wherein said second gateway is a first-hop away from said sender
18 host.

1 28. The system as recited in claim 26, wherein said first TCP packet of data is sent
2 from said sender host to said receiver host through said first gateway a consecutive
3 number of times, wherein said sender host fails to receive an acknowledgment of
4 received data from said receiver host each time said first TCP packet of data is sent to
5 said receiver host.

1 29. The system as recited in claim 26, wherein said routing table in said sender host
2 is located in a network layer of a TCP/IP protocol suite.

1 30. The system as recited in claim 29, wherein said network layer is an IP layer.

1 31. The system as recited in claim 26, wherein said sender host further comprises:
2 circuitry for establishing a new communication using said first gateway by said
3 application of said sender host.

1 32. The system as recited in claim 26, wherein said sender host further comprises:
2 circuitry for establishing a new communication using said first gateway by a new
3 application of said sender host

1 33. The system as recited in claim 31, wherein said sender host further comprises:
2 circuitry for sending an ARP request to said first gateway by said sender host.

1 34. The system as recited in claim 33, wherein said ARP request is sent from said
2 sender host to said first gateway a consecutive number of times.

1 35. The system as recited in claim 33, wherein if said sender host receives a response
2 from said first gateway then said first gateway is operative, wherein if said sender host
3 receives a response from said first gateway then said application of said sender host
4 sends a second TCP packet of data using said first gateway if said new communication
5 is a TCP communication.

1 36. The system as recited in claim 33, wherein if said sender host receives a response
2 from said first gateway then said first gateway is operative, wherein if said sender host
3 receives a response from said first gateway then said application of said sender host
4 sends a non-TCP packet of data using said first gateway if said new communication is
5 a non-TCP communication.

1 37. The system as recited in claim 33, wherein if said sender host does not receive
2 a response from said first gateway then said first gateway is inoperative, wherein if said
3 sender host does not receives a response from said first gateway then said sender host
4 further comprises:

5 circuitry for selecting an alternative path through an alterative gateway than said
6 first gateway in said routing table in said sender host.

1 38. The system as recited in claim 37, where said application of said sender host
2 sends a non-TCP packet of data using said alternative gateway if said new
3 communication is a non-TCP communication.

4 39. The system as recited in claim 37, where said application of said sender host
5 sends a second TCP packet of data using said alternative gateway if said new
6 communication is a TCP communication.

1 40. The system as recited in claim 37, wherein said sender host further comprises:
2 circuitry for marking all routes that use said first gateway to a lower priority level
3 from an original priority level in said routing table in said sender host.

4 41. The system as recited in claim 40, wherein said all routes that use said first
5 gateway return to their original priority level after a duration of time.

1 42. The system as recited in claim 32, wherein said sender host further comprises:
2 circuitry for sending an ARP request to said first gateway by said sender host.

1 43. The system as recited in claim 42, wherein said ARP request is sent from said
2 sender host to said first gateway a consecutive number of times.

1 44. The system as recited in claim 42, wherein if said sender host receives a response
2 from said first gateway then said first gateway is operative, wherein if said sender host

3 receives a response from said first gateway then said new application of said sender host
4 sends a second TCP packet of data using said first gateway if said new communication
5 is a TCP communication.

1 45. The system as recited in claim 42, wherein if said sender host receives a response
2 from said first gateway then said first gateway is operative, wherein if said sender host
3 receives a response from said first gateway then said new application of said sender host
4 sends a non-TCP packet of data using said first gateway if said new communication is
5 a non-TCP communication.

1 46. The system as recited in claim 42, wherein if said sender host does not receive
2 a response from said first gateway then said first gateway is inoperative, wherein if said
3 sender host does not receives a response from said first gateway then said sender host
4 further comprises:
5

1 circuitry for selecting an alternative path through an alterative gateway than said
2 first gateway in said routing table in said sender host.
3
4
5

1 47. The system as recited in claim 46, where said new application of said sender host
2 sends a non-TCP packet of data using said alternative gateway if said new
3 communication is a non-TCP communication.
4
5

1 48. The system as recited in claim 46, where said new application of said sender host
2 sends a second TCP packet of data using said alternative gateway if said new
3 communication is a TCP communication.

1 49. The system as recited in claim 46, wherein said sender host further comprises:
2 circuitry for marking all routes that use said first gateway to a lower priority level
3 from an original priority level in said routing table in said sender host.

1 50. The system as recited in claim 49, wherein said all routes that use said first
2 gateway return to their original priority level after a duration of time.

1 51. A computer program product having a computer readable medium having
2 computer program logic recorded thereon for selecting an alternative path upon detection
3 of a dead gateway, comprising:

4 programming operable for sending a first TCP packet of data from an application
5 of a sender host to a receiver host through a first gateway, wherein said sender host fails
6 to receive an acknowledgment of received data from said receiver host;

7 programming operable for deleting an ARP entry associated with said first
8 gateway in said sender host; and

9 programming operable for selecting an alternative path to send said first TCP
10 packet of data from said sender host to said receiver host through a second gateway in
11 a routing table in said sender host.

12 52. The computer program product as recited in claim 51, wherein said first gateway
13 is a first-hop away from said sender host, wherein said second gateway is a first-hop
14 away from said sender host.

15 53. The computer program product as recited in claim 51, wherein said first TCP
16 packet of data is sent from said sender host to said receiver host through said first
17 gateway a consecutive number of times, wherein said sender host fails to receive an
18 acknowledgment of received data from said receiver host each time said first TCP packet
19 of data is sent to said receiver host.

20 54. The computer program product as recited in claim 51, wherein said routing table
21 in said sender host is located in a network layer of a TCP/IP protocol suite.

22 55. The computer program product as recited in claim 54, wherein said network layer
23 is an IP layer.

24 56. The computer program product as recited in claim 51 further comprises:

programming operable for establishing a new communication using said first gateway by said application of said sender host.

57. The computer program product as recited in claim 51 further comprises:
programming operable for establishing a new communication using said first gateway by a new application of said sender host

58. The computer program product as recited in claim 56 further comprises:
programming operable for sending an ARP request to said first gateway by said sender host.

59. The computer program product as recited in claim 58, wherein said ARP request is sent from said sender host to said first gateway a consecutive number of times.

60. The computer program product as recited in claim 58, wherein if said sender host receives a response from said first gateway then said first gateway is operative, wherein if said sender host receives a response from said first gateway then said application of said sender host sends a second TCP packet of data using said first gateway if said new communication is a TCP communication.

61. The computer program product as recited in claim 58, wherein if said sender host receives a response from said first gateway then said first gateway is operative, wherein if said sender host receives a response from said first gateway then said application of said sender host sends a non-TCP packet of data using said first gateway if said new communication is a non-TCP communication.

62. The computer program product as recited in claim 58, wherein if said sender host does not receive a response from said first gateway then said first gateway is inoperative, wherein if said sender host does not receives a response from said first gateway then the computer program product further comprises:

5
6

- 1
- 2
- 3

- 1
- 2
- 3

1
2
3
4

1000

1
2

34

12

1 69. The computer program product as recited in claim 67, wherein if said sender host
2 receives a response from said first gateway then said first gateway is operative, wherein
3 if said sender host receives a response from said first gateway then said new application
4 of said sender host sends a second TCP packet of data using said first gateway if said
5 new communication is a TCP communication.

1 70. The computer program product as recited in claim 67, wherein if said sender host
2 receives a response from said first gateway then said first gateway is operative, wherein
3 if said sender host receives a response from said first gateway then said new application
4 of said sender host sends a non-TCP packet of data using said first gateway if said new
5 communication is a non-TCP communication.

1 71. The computer program product as recited in claim 67, wherein if said sender host
2 does not receive a response from said first gateway then said first gateway is inoperative,
3 wherein if said sender host does not receives a response from said first gateway then the
4 computer program product further comprises:

5 programming operable for selecting an alternative path through an alterative
6 gateway than said first gateway in said routing table in said sender host.

1 72. The computer program product as recited in claim 71, where said new application
2 of said sender host sends a non-TCP packet of data using said alternative gateway if said
3 new communication is a non-TCP communication.

1 73. The computer program product as recited in claim 71, where said new application
2 of said sender host sends a second TCP packet of data using said alternative gateway if
3 said new communication is a TCP communication.

1 74. The computer program product as recited in claim 71 further comprises:
2 programming operable for marking all routes that use said first gateway to a
3 lower priority level from an original priority level in said routing table in said sender
4 host.

1 75. The computer program product as recited in claim 74, wherein said all routes that
2 use said first gateway return to their original priority level after a duration of time.

DETECTION OF DEAD GATEWAYS**ABSTRACT OF THE INVENTION**

5 A method, computer program product and system for detecting a first-hop dead gateway. In one embodiment, a method comprises the step of sending a TCP packet of data from an application of a sender host to a receiver host through a first gateway, where the first gateway is a first-hop away from the sender host. The method further comprises the step of TCP failing to receive an acknowledgment of received data from the receiver host. The method further comprises the step of deleting an ARP entry associated with the first gateway in the sender host. The method further comprises the step of establishing a new communication using the first gateway by the application or new application of the sender host. The method further comprises the step of sending an ARP request to the first gateway by the sender host. If the sender host receives a response from the first gateway, then the method further comprises the step of sending a TCP or non-TCP packet of data using the first gateway if the new communication was a TCP or non-TCP communication, respectively. If the sender host does not receive a response from the first gateway, then the method further comprises the step of selecting an alternative path through an alternative first-hop gateway in the routing table of the sender host. The application or new application of the sender host then sends the TCP or non-TCP packet of data using the alternative gateway if the new communication was a TCP or non-TCP communication, respectively.


25 ::ODMA\PCDOCS\AUSTIN_1\146362\1
1162:7047-P366US

Figure 1

100

FTP	Telnet	HTTP	SNMP	RIP	150	
TCP		UDP		140		
IP and ICMP				RIP	ARP	130
Ethernet, Token Ring, IEEE 802.3					120	
ATM, CSMA/CD, Packet Switching					110	

Figure 2




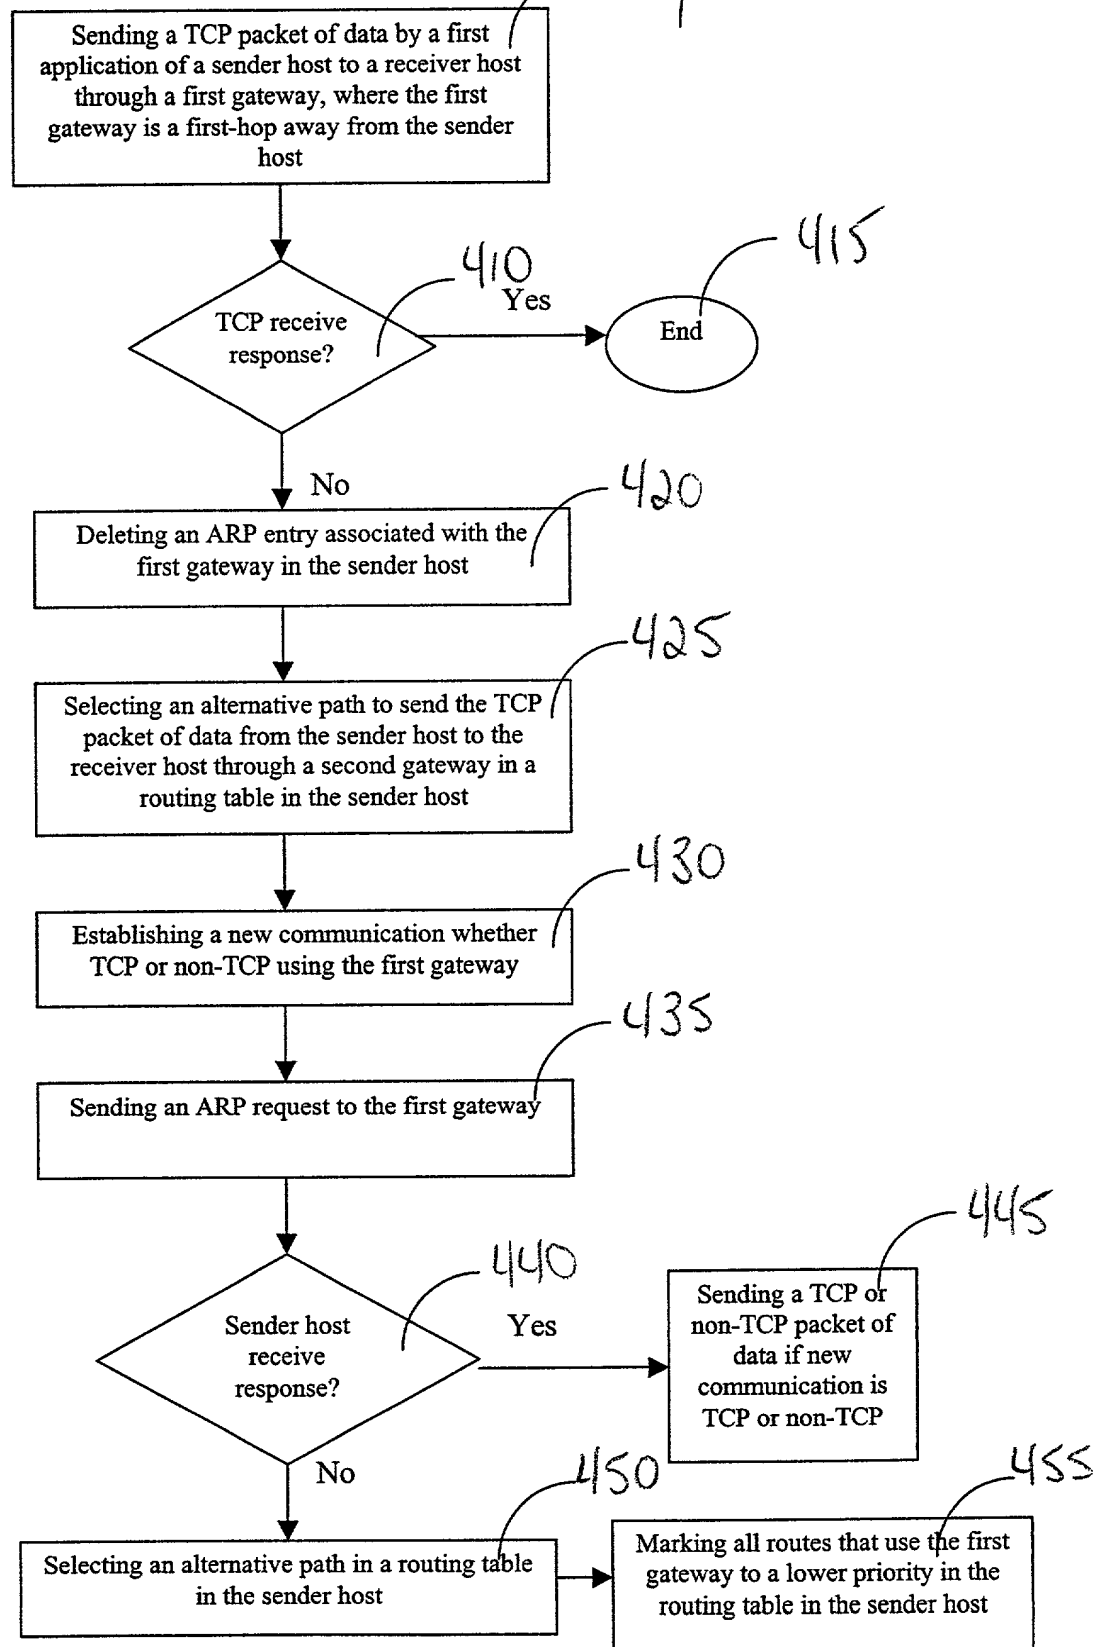


FIG. 3

Figure 4



**DECLARATION AND POWER OF ATTORNEY FOR
PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

DETECTION OF DEAD GATEWAYS

the specification of which (check one)

- ☒ is attached hereto.
- ☐ was filed on _____
as Application Serial No. _____
and was amended on _____

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s):

Priority Claimed

(Number)

(Country)

(Day/Month/Year)

☐ Yes ☐ No

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial #)

(Filing Date)

(Status)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

John W. Henderson, Jr., Reg. No. 26,907; James H. Barksdale, Jr., Reg. No. 24,091; Thomas E. Tyson, Reg. No. 28,543; Robert M. Carwell, Reg. No. 28,499; Jeffrey S. LaBaw, Reg. No. 31,633; Douglas H. Lefevre, Reg. No. 26,193; Casimer K. Salys, Reg. No. 28,900; David A. Mims, Jr., Reg. No. 32,708; Mark E. McBurney, Reg. No. 33,114; Anthony V. S. England, Reg. No. 35,129; Volel Emile, Reg. No. 39,969; Christopher A. Hughes, Reg. No. 26,914; Edward A. Pennington, Reg. No. 32,588; John E. Hoel, Reg. No. 26,279; Joseph C. Redmond, Jr., Reg. No. 18,753; Leslie A. Van Leeuwen, Reg. No. 42,196; Marilyn S. Dawkins, Reg. No. 31,140; Kelly K. Kordzik, Reg. No. 36,571; Barry S. Newberger, Reg. No. 41,527; and Robert A. Voigt, Jr., Reg. No. P47,159.

Send correspondence to: Kelly K. Kordzik, 100 Congress Avenue, Suite 800, Austin, Texas 78701, and direct all telephone calls to Kelly K. Kordzik at (512) 370-2851.

FULL NAME OF FIRST OR SOLE INVENTOR: **DEANNA LYNN QUIGG BROWN**

INVENTOR'S SIGNATURE: Deanna L. Brown

DATE: 8/29/00

RESIDENCE: **20824 Derby Day Avenue**
Pflugerville, Travis County, Texas 78660

CITIZENSHIP: **U.S.A.**

POST OFFICE ADDRESS: **(Same as Residence)**

FULL NAME OF SECOND INVENTOR: **VINIT JAIN**

INVENTOR'S SIGNATURE: Vinit Jain

DATE: 8/29/00

RESIDENCE: **16007 Rustic Lane**
Austin, Williamson County, Texas 78717

CITIZENSHIP: **INDIA**

POST OFFICE ADDRESS: **(Same as Residence)**

FULL NAME OF THIRD INVENTOR: **SATYA PRAKESH SHARMA**

INVENTOR'S SIGNATURE: Satya P. Sharma

DATE: 9/11/00

RESIDENCE: **1210 Robin Trail**
Round Rock, Williamson County, Texas 78681

CITIZENSHIP: **U.S.A.**

POST OFFICE ADDRESS: **(Same as Residence)**